codecentric

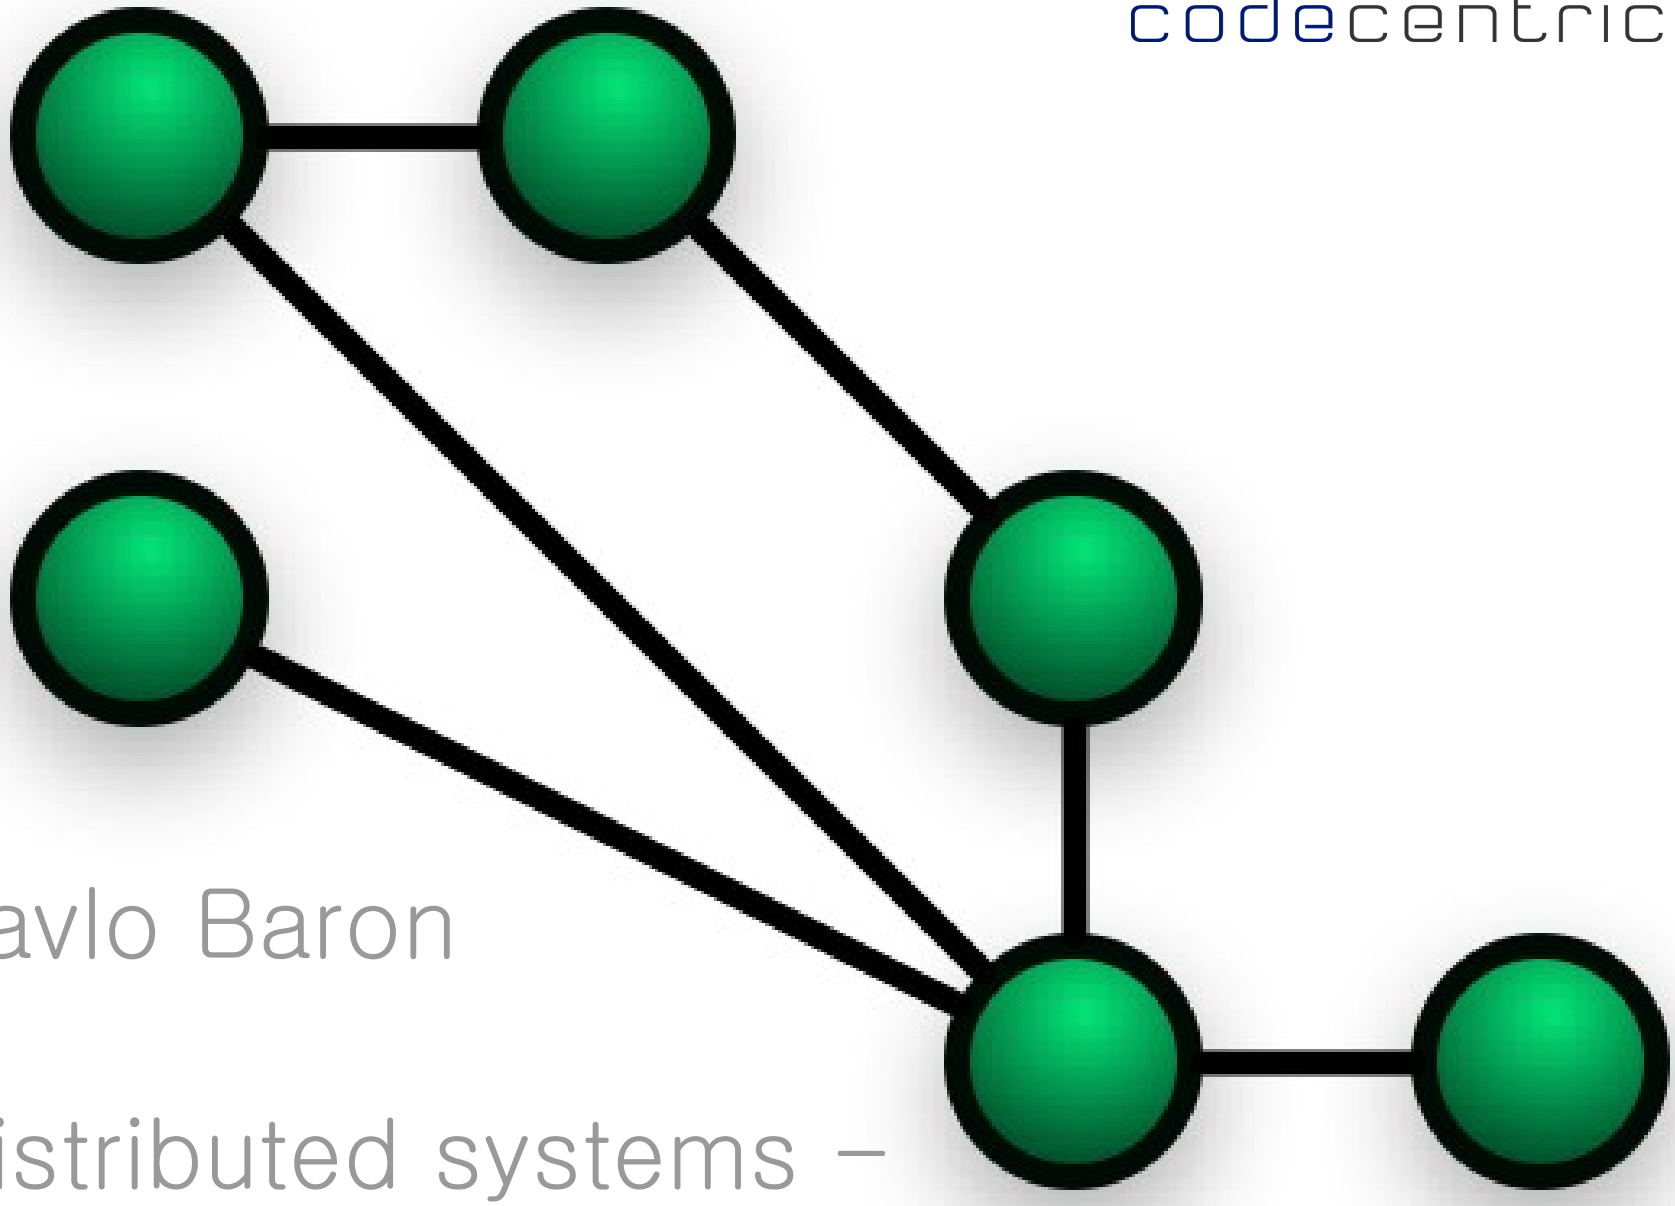

Pavlo Baron

Distributed systems – playfully illustrated

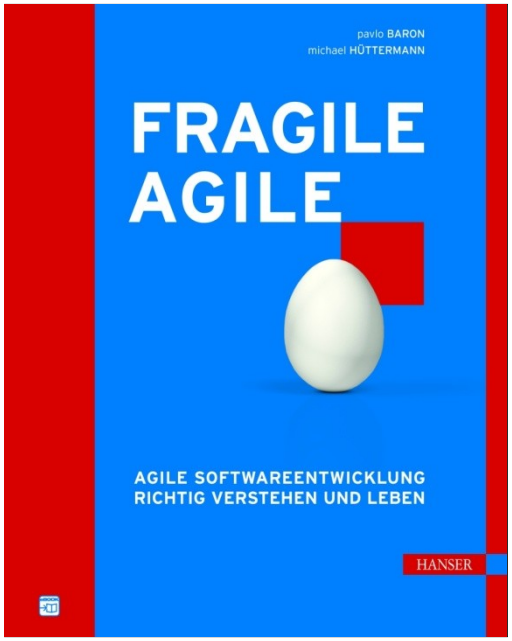

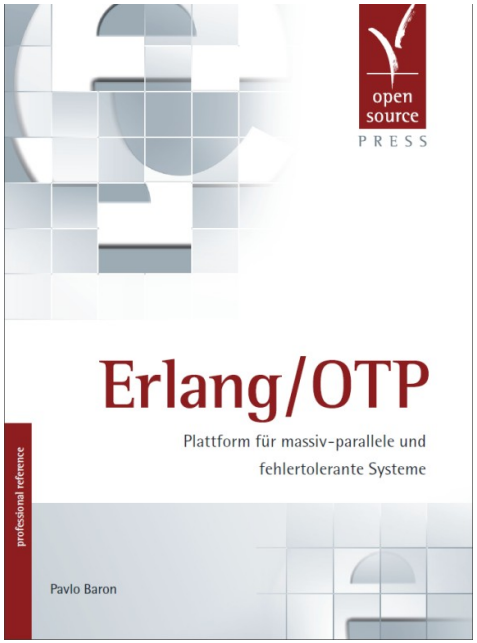

Geek's
Guide
To The Working Life

Pavlo Baron

pavlo.baron@codecentric.de
@pavlobaron

Distributed Systems – attributes & difficulties

Concurrency
Varying clocks
Independent failures
Different hardware & OSs
Easy localization
Softened consistency
Imperfect security

Distributed Systems – considerations & toolbox

Availability vs. consistency, latency
Asynchronous message passing
Replication
Caching & holding back
Equality of nodes
Optimism
No global state, no single points of error
Fault tolerance
Hashing, no sequencing
Concept recycling


CIRCUS AMOK

Example: clocks

Computer clocks are inconsistent.
Time is relative.
Global clocks are weak.

Logical / vector clocks allow to keep track of changes without a need to synchronize time

Example: consistent hashing

It is necessary to localize data on changing infrastructure with ~O(1) efficiency and minimal reorganization

Using ring hashing with fixed rules allows to hit the moving target

Example: gossip

Notification of data and infrastructure changes must happen asynchronously. Notifications occure in unpredictable order and repetition.

Nodes gossip in order to share these changes, information is sent piggyback

Example: hinted handoff

Nodes and whole network segments can fail.

In order to prevent data lost another nodes can buffer foreign data and hand it off later to reanimated nodes

Example: quorum

Softened consistency still needs assurance that data has been reliably read and written

With „W > 0.5 V” und “R + W > V” rules consistency can be reliably softened

Example: election

When all nodes are equal, there is still need to choose a coordinator for some processes

Election allows to decide which node plays a certain role (ring, bully)

Example: failure detection

Reliable / accurate failure detectors are impossible in an asynchronous system

Optimistic failure detection allows to suspect a node of having failed, exclude it, while still hoping it will come back

Example: partion tolerance / no split brain

When the overall network suddenly gets partitioned, nodes must not think they are all alone while still doing their job

Pessimistic tolerance eventually means data lost, optimistic means weaker consistency

Thank you

Most images originate from
istockphoto.com

except few ones taken
from Wikipedia and product pages
or generated through public
online generators